

Use of a Job Ticket as a Generic XML Database

Technical Field

The technical field is the integration and control of services in a networked environment.

Background

Services may be provided by one or more operating units in a computer-based network. Users of the network may generate specific tasks and send the tasks into the network to be assigned to one of the operating units. For example, a user at a computer terminal may generate a printing order using a printer driver installed on the terminal. The printer driver is used to control the printing request. In another example, a user at a computer terminal may generate a printing order and send the printing order into a computer network so that the printing order is completed by a printing service. The printing order may be related to a company brochure. The printing order may contain unique requirements such as paper type, font size, layout, graphics, color, and other requirements. The user may specify that a specific printing service, such as Kinkos, prepare the company brochure. Alternatively, the computer network may include programs that suggest printing services to the user.

To control the printing job, the user's computer terminal may generate a job ticket. The job ticket includes the requirements, such as the requirements listed above, and an identification of the specific job that allows the job status to be tracked through the computer network.

Use of the job ticket allows printing and similar services to be allocated to those resources (i.e., the operating units) that are best suited to completing the services. Unfortunately, current computer systems do not allow access to the wide variety of services existing in networked computer systems, such as the Internet. In addition, current systems require users to have some knowledge of the existing resources, and may require users to include applicable programming to communicate with the services. Furthermore, current systems do not allow a job request to be split among several processors. As a result, completion of the job request may take longer than necessary, and may not be completed in the most efficient, lowest-cost manner.

Summary

To overcome these and other problems related to use of a job ticket, a method and an apparatus allow a client to manage job attributes and processes using an electronic service center. The service center includes a job ticket service that allows access and modification of a job ticket

by multiple users on a network. The method and apparatus use a network-accessible job ticket to relate to a specific job or content. The job ticket may be an object, such as an XML object, comprising routines and data. The content may be stored on the network and may be accessed by multiple job tickets. Storage and management of the job ticket are transparent to the user. The job ticket is stored in a common location in the network. The job ticket remains in the same location in the network, and users access only that portion of the job ticket required to complete a designated process. Security measures may be added to limit access to those users designated as being allowed to access the job ticket and the job file. The job ticket may include a service ID that relates the job ticket back to the originating job ticket service. In this way, a user who acquires all or part of the job ticket can refer back to the originating job ticket service (and the original, or as-modified, job ticket) to verify any changes and to ensure that the job ticket being accessed is up-to-date. The job ticket also includes a job ID to refer the job ticket to a specific job.

The service center is coupled through a communications network to a front end service. The front end service allows a user to generate a service or job request. The communications network may be the Internet, or a local area network, for example.

The service center includes a service bus, to which are coupled a job store, the job ticket service, and a work flow controller. Also coupled to the service center are one or more processors that may be controlled to complete processes and tasks defined in the job tickets.

The job ticket service may generate and store the job tickets. Job content (e.g., a PDF file) is stored in the job store. With this structure, the user does not have to manage storage of the job content or to know which job store holds the job content. The job ticket service controls access to the job tickets, and, through the use of the job tickets, also controls access to job content in the job store, or elsewhere in the network. The job ticket service may create a reference to the job ticket, and may use the reference to control access to the job ticket.

The use of job tickets allows clients to define databases, and to store data through the job ticket service. The databases may be used to hold contact lists, addresses, and other personal data. The databases may also be used to store any other generic data. The databases could then be used in conjunction with a variety of e-services provided by the processors. For example, an e-mail processor that provides e-mail services may be used in conjunction with a personal contact list to send e-mail messages, transfer electronic files, or to establish a chat room. The e-mail

processor may access the contact list at predefined intervals to send e-mail messages to a select group of e-mail addressees. Furthermore, because the service center provides a single portal to processors that are coupled to the communications network, the client need not have any knowledge of the database structure, or the processing requirements of the processors.

In an embodiment, the job ticket may be used as a generic database in conjunction with an e-mail service. A client may have established a list of e-mail contacts. The contacts database may then be stored in the job store. A corresponding job ticket may be stored at the job ticket service. The job ticket includes control data needed to send and receive e-mail through the service center. In an embodiment, the generic database may be an XML or HTML database, or any other markup language database. The database may also be a structured query language (SQL) database.

The use of the job ticket also allows for parsing, searching and updating the contacts database. For example, the client may desire to search the contacts database for phone numbers by name. This search functionality is included in the job ticket, and allows the job ticket service to provide the client with a list of phone numbers for all entries in the contacts database by name.

Description of the Drawings

The detailed description will refer to the following figures in which like numeral refer to like items, and in which:

Figure 1 is a block diagram showing a prior art use of a job ticket;

Figure 2 is a tree diagram showing the processes in an example job ticket;

Figure 3 is a block diagram of a digital image work flow network;

Figure 4 is a block diagram of a service center used with the network of Figure 3;

Figures 5A-5D illustrate an exemplary job ticket;

Figure 6 is a diagram of functions controlled by a job ticket service;

Figure 7 is a diagram showing access functions controlled by the job ticket service;

Figure 8 is a block diagram illustrating additional control features of the job ticket service;

Figure 9 is a flow chart illustrating one of the processes controlled by the job ticket service;

and

Figures 10-16 are flow charts showing sub-processes in the overall process illustrated in Figure 9.

Detailed Description

Figure 1 is a block diagram showing a prior art application of a job ticket service. Job tickets are often associated with a printing standard, the job definition format (JDF). The JDF is described in detail in JDF Specification Draft Spiral 4.0, available at www.hp_opensource.com, which is hereby incorporated by reference. In Figure 1, a user 1 generates a job request and sends the job request through a portal 4 to a processor 5. The job request may include a job ticket data file 2 and a content file 3. The user 1 may be a computer terminal in a networked computer system and the processor 5 may be a networked printer. The job request may involve printing a document. The document may be represented by the content 3, which is a digital representation of text and images to be printed. The intended format of the printed document may be described in the job ticket file 2, which is simply a digital file that specifies how the printer is to print the document. For example, the job ticket file 2 may require that the document be printed on back-to-back pages.

In a specific application, the functions of the job ticket file 2 may be carried out by a printer driver. The printer driver encodes control data related to printing the document, and sends the control data and the content 3 to the printer (i.e., the processor 5). The printer accesses the control data and the content 3 to print the document.

While the application shown in Figure 1 works well to print a document, the application has many drawbacks. In particular, if multiple processors are involved in producing the document, each such processor will require access to the job ticket file 2. This access brings problems related to security, modification control and workflow control. For example, each processor requiring access to the job ticket file 2 may have to wait on processing until a prior processor has completed use of the job ticket file 2. Thus, the prior art application may result in unwanted delays in completing the job request.

Prior art applications of job ticket services also suffer because the user may not know anything about the processors, including capabilities and availabilities of the processors, or even if the processors exist. Thus, the user may not know which portal to use to connect to a specific processor.

These and other problems are solved by a method and an apparatus that controls access to a job ticket and associated content through use of a job ticket service. The job ticket service

includes mechanisms that arbitrate access to the job ticket among multiple users of the job ticket, limit access to the job ticket by incorporating security features, and ensure modifications made by one processor or user are reflected in the job ticket and the content. In effect, the apparatus includes a generic database that couples input data from clients as job requests with output services such as processors that perform tasks or processes to complete the job requests. The database may have has the features of a generic XML database in that it is extensible, and in that the clients need not have any knowledge of the individual processes to be performed, or the internal programming requirements of the processors. Thus, the clients may submit job requests to a service center that will ensure that an appropriate processor or processors are assigned to complete the job request. Other database formats may also be used.

Before describing the apparatus and method in detail, a review of a job ticket is provided. Figure 2 is a node-tree diagram (or simply a node tree) 10 that illustrates processes defined in a job ticket for printing a brochure. The brochure may be printed on a commercial press, and may use digital content to generate plates for printing the brochure. Within the node tree 10, the nodes specify a product, process, or group of processes. Each node may modify, consume or create resources. Each node may contain further nested nodes, or sub-nodes. The arrangement of nodes and sub-nodes may be likened to a tree, and each node and sub-node may be referred to as a branch. A brochure node 11 defines the features and parameters of the brochure. A cover node 12 defines the parameters for producing the brochure cover. Inside pages node 13 includes the parameters to produce the inside pages. The inside pages node 13 is shown with several sub-nodes, including a sub-node 14 for digital plate making. The digital plate making sub-node 14 itself includes two additional sub-nodes, a ripping sub-node 16 and a plate making sub-node 18.

Each of the nodes and sub-nodes shown in Figure 2 has associated with it input resources and at least one output resource. A resource may be described by parameters or logical entities. The resource may be a physical entity such as a component, a handling resource, or a consumable. A component resource may be the output of a node or sub-node, such as a printed sheets. A handling resource is used during a process, but is not consumed by the process. A consumable resource may be partly or wholly consumed by the process. Examples of consumable resources include inks, plates, and glue. Other resources may be a digital file or representation of a physical object. For example, the ripping sub-node 16 may include as input resources a run list, media, RIP

parameters, and layout. The run list resource describes the pages, including the files in which the pages occur, and which pages are to be used. The media resource describes the media that will be used to make plates, and is needed to describe the dimensions of the media. The RIP parameters resource describes all device-specific parameters of the ripping process. The layout resource describes placement of source pages onto the plates, and eventually onto press sheets. As an output resource, the ripping sub-node 16 may provide ripped flats. Other resources include parameter resources, which define the details of processes, as well as other non-physical computer files used by a process.

The node tree 10 shown in Figure 2 is intended to apply to printing a document. However, node-tree diagrams may be used to represent job tickets for other services besides printing. For example, a job ticket may be used for data processing, image processing, creating and maintaining a database, electronic publishing, e-mail, and various e-commerce services. Moreover, the job ticket may be used to allow different e-commerce services to interact with each other.

Figure 3 is a block diagram of a digital imaging work flow (DIW) network 20 that incorporates a service center and a job ticket service to control tasks submitted by clients. The service center may operate as a single portal through which the clients connect to one or more e-services including e-mail, e-commerce and online shopping, e-printing, and data services, including database searching and database construction, population and maintenance. Using a single portal, such as the service center, allows the clients to select from a wide variety of e-services, such as those noted above, without requiring the clients to have any prior knowledge of the e-services.

The service center may include components that receive information in the form of job requests, and using the information, create a job ticket that specifies tasks and resources. The job ticket may be stored in a job ticket service, and notices may be posted to indicate when a job ticket is available. Processors coupled to the service center may bid on completion of the job ticket, and the service center may include a bidding service that evaluates bids. The service center may select one or more processors to assign to the job ticket based on client-supplied criteria, or based on a set of standard criteria, including industry standard criteria. The service center may provide mechanisms to control access to the job ticket, or to portions (branches) of the job ticket. The mechanisms include branch locking, and authorization and authentication servers that use public key encryption, or similar processes.

1 The service center may include hardware components such as servers, computers, central
2 processing units, communications interfaces, and memory devices to provide the processing
3 capability and data storage required to carry out the above-described functions.

4 The DIW network 20 includes a front end service 30 that allows a client 31 to generate
5 and submit a service or job request. In an embodiment, the front end service 30 may be an Internet
6 web browser. Alternatively, the front end service 30 may be a web application or a port monitor.
7 The job request may contain detailed information about how the job is to be executed, and may
8 be formatted according to the job definition format standard. Alternatively, the job request may
9 include only basic information, which will be used by another component to finalize the job
10 definition, or work flow. Finally, the job request may include the content, or job, that is to be
11 processed. The content could be one or more digital files, text files, and other files. The front end
12 service 30 is coupled to a communications network 35, which may be the Internet or a local area
13 network, for example. Coupled to the communications network 35 is a service center 40 that links
14 one or more processors 80_i to the communications network 35. Each of the processors 80_i may
15 include a cache 81_i that may be used to store information related to a job request, including
16 information related to job tickets. In an embodiment, the service center 40 may be an Internet web
17 site that includes data storage and control functions. In another embodiment, the service center 40
18 is a node in a local area network.

19 The service center 40 allows a broad spectrum of communications between entities coupled
20 to the service center 40. In particular, the service center 40 allows different e-services to interact
21 programmatically with one another using specific protocols and generic protocols (e.g., TCP/IP).
22 This programmatic interaction allows different services and processes that are coupled to the
23 network to exchange data and files, and to modify the data and files. The programmatic interaction
24 may be completed by use of a remote procedure call (RPC) between entities coupled to the service
25 center 40. Other methods for providing the programmatic interaction include CORBA, UDDI, and
26 e-speak.

27 Figure 4 is a diagram of the service center 40. The service center 40 includes a service bus
28 41 in communication with the communications network 35 and the processors 80_i of Figure 3.
29 Coupled to the service bus 41 is a job store 50, a job ticket service 60, a work flow controller 70,
30 an optional bidding service 90, an authorization server 92 and an authentication server 94. The job

1 store 50 may store one or more job content files 51_i. The job ticket service 60 may control one
 2 or more job tickets 61_i. The work flow controller 70 may use one or more agents 71_i to control
 3 processes on the service bus 41.

4 The job store 50, job ticket service 60 and work flow controller 70 function to accept
 5 information from the clients 31, and to use the information to control the actions of the processors
 6 80_i. The processors 80_i performs specific tasks or processes as determined by the service center
 7 40.

8 The job store 50 may be a node on the service bus 41, and may include programming to
 9 allow the job store 50 to carry out its functions. The job store 50 may be used to store the content
 10 51, which may be in the form of one or more large files. In the context of printing a document using
 11 a service or process coupled to the service bus 41, the job store 50 may store the document
 12 content in one or more PDF files, for example. The content 51 may include graphics and text. The
 13 content 51 for a specific document may include several files. For example, a brochure may have
 14 a separate file for the cover and another file for the inside pages. Text for the inside pages may be
 15 in one file and images in yet another file. The content 51 may also include links to other resources
 16 or entities on the service bus 41. The job store 50 provides for mass storage of the content 51, so
 17 that a user (client 31 or processor 80) does not have to provide the mass storage required for the
 18 job content 51. By using the mass storage capabilities of the job store 50, the content 51 may be
 19 made to persist in the network 20, and may be made accessible to users at any time. That is, the
 20 job store 50 may store the content 51 for an extended time. The job store 50 also manages and
 21 controls the content so that the user (client 31 or processor 80) does not have to manage the
 22 content 51. Management functions include maintaining configuration or version control of the
 23 content 51, controlling access to the content 51, and maintaining the content 51 in storage.

24 The job ticket service 60 holds job tickets 61. The job ticket service 60 controls access
 25 to and may manage configuration of the job tickets 61. For example, the job ticket service 60 may
 26 allow users (clients 31_i and processors 80_i) to access a portion or branch of a job ticket 61 rather
 27 than passing the job ticket 61 among multiple users. Access to the job ticket portion may be
 28 effectuated by use of an application programming interface, a scriptable interface, or a similar
 29 feature. As noted above, the job ticket 61 does not include the content 51 (e.g., the graphical and
 30 text files of a document), but the job ticket 61 relates to content 51 (e.g., a PDF file) stored in the

1 job store 50. The user does not have to manage storage of the job content or to know which job
2 store 50 holds the job content. The job ticket service 60 instead passes a reference in the job
3 ticket 61. This allows multiple clients 31 and processors 80_i to access the content 51.
4 Furthermore, the content 51 may relate to more than one job ticket 61. The job ticket service 60,
5 and its interrelationships with other entities coupled to the service bus 41, will be described later
6 in detail.

7 Some job tickets 61 may be accessed by multiple processors 80, in either serial,
8 overlapping, or simultaneous fashion. The multiple access processing could result in problems with
9 use of the job ticket 61. For example, a first processor may acquire the job ticket 61 (or a portion
10 or branch thereof), and perform a process specified in the work flow, which may modify the
11 branch. Such modification may be to indicate a branch as complete, use up input resources, or
12 create new output resources, for example. A second processor could attempt to acquire the
13 branch, but might not “know” that the first processor had modified the branch. Alternatively, if two
14 processors compete for the same branch, a deadlock situation might occur.

15 One solution to the above problems may be to lock the job ticket 61 whenever a processor
16 80 acquires the job ticket 61. Unfortunately, locking the job ticket 61 may prevent concurrent or
17 parallel processing and may slow down completion of the job request 32.

18 The job ticket service 60 shown in Figure 4 overcomes these and other problems by having
19 the capability to lock the job ticket 61 at the branch level. The branch locking may be
20 accomplished by one of several methods. The work flow controller 70 may assign one or more
21 specific processors 80_i to perform the tasks identified with the branch to be locked. Where only
22 one processor 80 is authorized access to the branch, branch locking may not be required. Where
23 more than one processor 80 is authorized access to the same branch, the job ticket service 60 may
24 lock the branch when one of the authorized processors 80 actually acquire the branch.

25 If the work flow controller 70 has not assigned processors 80_i to branches (i.e., any
26 processor 80 may access a branch at any time), the job ticket service 60 may lock the branch when
27 a processor 80 acquires the branch.

28 The job ticket service 60 may lock the branches by setting a lock/unlock flag for each
29 branch. Processors 80_i accessing the job ticket 61 may then review the lock/unlock flag status to
30 determine if the branch may be accessed. In some circumstances, the job ticket service 60 may

allow access only to those branches that are unlocked. A processor 80 that has completed a task defined by the branch may need to have the branch unlocked in order to modify the branch.

The work flow controller 70 may be used to create the job tickets 61 that are stored in the job ticket service 60. The work flow controller 70 may review the job requests 32 submitted by the clients 31, and may then use a job ticket template to prepare the job ticket 61. The work flow controller 70 may then send the job ticket 61 to the job ticket service 60 for storage and processing.

The work flow controller 70 also controls completion of tasks among the processors 80_i. In an embodiment, the work flow controller 70 determines which of the processors 80 have the necessary and available resources to begin the processes listed in a specific job ticket 61. The work flow controller 70 then designates the appropriate processors 80_i to complete the tasks referenced by the job ticket 61. For example, if a job ticket 61₁ requires color printing, the work flow controller 70 may determine that only processor 80₃ is a color printer with the capacity to begin the job specified in the job ticket 61₁. This embodiment in which the work flow controller 70 determines which processors 80_i to assign to a specific job ticket 61 may be especially appropriate when the network 35 is a local area network and all processors 80_i are directly coupled to the local area network 35.

Alternatively, the work flow controller 70 may receive bid information from Internet-connected processors 80_i and may use the bid information to select the processors 80_i to complete the job request 32.

The work flow controller 70 may also be used to designate the various nodes, input and output resources, and other features of the node tree used to complete the job request. That is, the work flow controller 70 may be used to create a construct, or work flow, such as the node tree 10 shown in Figure 2. To accomplish these tasks, the work flow controller 70 may include one or more agents 71_i that write a job definition file, based on control data contained in the job request 32. Alternatively, a separate management information system (not shown) may be used to create the nodes, and to control flow of tasks to the processors 80_i and other entities. In yet another embodiment, the job definitions may be written by the client 31 that originated the job request 32.

Referring again to the node tree 10 of Figure 2, many output resources of the individual nodes serve as input resources for other nodes. These other nodes may not be able to begin

executing until all input resources are complete and available, which means that the nodes may need to execute in a well-defined sequence. For example, a process for making plates will produce press plates as an output resource that is required by a printing process. In the hierarchical organization of the node tree 10, nodes that occur higher in the node tree 10 represent higher-level, more abstract operations, while lower order nodes represent more detailed, specific processes. Moreover, nodes near the top of the node tree 10 may represent only intent regarding the components or assemblies that comprise the product, and lower level nodes provided the detailed instructions to a processor 80 to perform a specific process.

Because two node trees may not be similar, the work flow controller 70 may determine processes to be completed, the order in which the processes are completed, and the processors 80_i that are to complete the processes. The work flow controller 70 may use the agents 71 to determine an actual work flow, considering factors such as control abilities of the processors 80_i that complete the processes, transport distances between processors 80_i, load capabilities of the processors, and time constraints in the job request, for example. The agents 71 may define the overall process using serial processing, which involves subsequent production and consumption of resources by the processors 80_i, overlapping processing, which involves simultaneous consumption and production of resources by more than one processor 80_i, parallel processing, which involves sharing resources among processors 80, and iterative processing, which involves a back and forth processing scheme to develop resources.

Returning to Figure 4, the processors 80_i may be used to provide data services to the clients 31. However, each processor 80 can have a unique interface, data format, query syntax, security protocol, and the like. For example, the processor 80₁ may be configured as a web server and so responds to hypertext transfer protocol (HTTP) request packets and generates responses formatted as hypertext markup language (HTML) web pages. In contrast other processors 80_i may use file transfer protocol (FTP) or other available transfer protocols, both public and private.

In a particular example, each of the processors 80₁ - 80₃ implements a database web server that hosts a "website" having a predefined structure and access syntax and protocol. In general each website is designed to provide information about its database content response to queries (job requests) communicated through network 20. Because each processor 80_i may be created and maintained by a separate organization, there is no uniform interface provided for

1 interacting with the various processors 80_i. For example, the databases may relate to customer lists
2 in an e-commerce application, and the processor 80₁ may enable searches by last name, phone
3 number, zip code, or other searchable fields whereas the processor 80₂ may enable searches only
4 by last name. Accordingly, and in the absence of the service center 40, the job requests 32 must
5 differ in content to evoke the desired response from each processor 80_i. Also, the processor 80₁
6 may require that a query be processed through a series of pages before a response is generated.
7 In contrast, the processor 80₂ may allow direct, one-page access to submit queries and generate
8 responses. Hence, the processor 80₁ will require a series of communication packets to be received
9 before the desired response is generated whereas the processor 80₂ is potentially accessed with
10 a single communication packet that encapsulates the desired query data.

11 The service center 40 serves as an interface between the client 31 and the various
12 processors 80_i, and allows the client to submit data queries (job requests) without having any prior
13 knowledge of the database being searched, and its associated search engines.

14 In determining which of the processors 80_i to assign to complete a particular job request,
15 the work flow controller 70 may poll the processors 80_i that are coupled to the service center 40.
16 As noted above, the processors 80_i may be coupled directly to the service bus 41, or may be
17 coupled indirectly through another communications bus, such as the Internet, for example. The
18 polling may occur whenever a job ticket 61 is created by the job ticket service 60. Alternatively,
19 the polling and corresponding information collection may occur on a periodic basis, and the work
20 flow controller 70 may store information related to the processors 80.

21 As an alternative to polling, processors 80_i coupled to the service center 40 may monitor
22 the job ticket service 60. The job ticket service 60 may periodically post, in a bulletin board
23 fashion, for example, notices for job tickets that are available for processing. The processors 80_i
24 may then submit a bid for the tasks and processes defined in the job ticket notice. The work flow
25 controller 70, or the separate, optional bidding service 90, may review the bids, and determine
26 which single processor 80 or combination of processors 80_i would be best suited to complete the
27 tasks and processes defined in the job ticket notice.

28 The service center 40 may include several features to provide security and to control access
29 to the job ticket 61. As discussed above, the job ticket service 60 may include a provision for
30 branch locking. In addition, servers may be used to authorize and authenticate a processor 80 and

1 maintain the authorization and authentication during completion of a job request 32. The
2 authentication server 92 receives authentication information from a processor 80 and the
3 authorization server 94 uses the information to check authorization functionality. The authorization
4 or access rights of the processor 80 may be carried as a part of the job ticket 61. The servers 92
5 and 94 may be hardware devices, but need not exist in the same hardware platform, and the
6 servers 92 and 94 need not be tightly coupled. Alternatively, the functions of the servers 92 and
7 94 may be performed in programming stored in one of the components of the service center 40,
8 such as the work flow controller 70, for example. Using the above-described features, the service
9 center 40 may provide trusted authentication information about the processor 80 to the
10 authorization server 94, and the authorization server 94 then performs its authority check functions.

11 The job ticket 61 may be signed with an industry standard public key encryption message
12 digest (MD) signature, and may be protected by a public key encryption system. Hence, any user
13 that has the public key may validate the job ticket 61 without having to communicate with the
14 authentication server 92. These features reduce communication between distributed server
15 applications. The features also allow the job ticket 61 to be passed from one processor 80 to
16 another processor 80, maintaining security, without communicating with the service center 40.

17 In an alternative embodiment, the job ticket 61 holds authentication/access data, allowing
18 controlled access within the service center 40 infrastructure. Resources may be protected by
19 passwords and other mechanisms. Access to the job ticket 61 may be similarly protected.
20 Furthermore, processors 80, with access authorization may have such access authorization invoked
21 by listing the processors in the job ticket. The listing may be effectuated by recording a network
22 address for the processors 80, for example. The network address may be incorporated in the bid
23 information recorded in the job ticket 61.

24 Although the above description refers to development by the work flow controller 70, other
25 components in the network 20 may be used to develop an overall work flow to complete the job
26 request 32. For example, the job ticket service 60 may be used to develop the overall work flow.

27 As discussed above, the bidding service 90 may be used to receive bid information from
28 processors 80, coupled to the service center 40. The processors 80, submit bids in response to
29 posting of job ticket notices at the service center 40. In an embodiment, the job ticket notice is a
30 separate object stored in the service center 40. In another embodiment, the job ticket 61 itself

1 serves the notice function. The work flow controller 70 may post the job ticket notices after receipt
2 of the job request 32. Whether the bidding service 90 or the work flow controller 70 receives the
3 bids, the bid evaluation and selection process may be the same.

4 The job ticket notice posted by the work flow controller 70 may include specific tasks or
5 processes (branches) that must be completed to complete the job request 32 (see Figure 7). A
6 simple job request 32 may have only one branch. More complex job requests 32, such as the job
7 request illustrated in Figure 2 (i.e., print a brochure) may have many branches. Furthermore, some
8 branches may be so interrelated that they can only be completed in a specific sequence, while other
9 branches can be completed in a parallel or an overlapping fashion. This interrelationship may often
10 be the result of one branch producing an output resource that is an input resource for one or more
11 other branches. The job ticket notice may include descriptions of specific branches and their
12 interrelationships in sufficient detail to allow the processors 80_i to bid for completion of the
13 branches. The job ticket notice may persist in the service center 40 for a specified time to allow
14 the processors 80_i to send bids. The time may be a set value (e.g., one hour) or may be based on
15 a completion deadline specified in the job request 32.

16 The bidding service 90 may select bids 91 from the processors 80_i based on set criteria.
17 For example, the job request 32 may specify minimum performance requirements (e.g., a maximum
18 cost and a completion deadline). The bidding service 90 may reject any bids that fail to satisfy the
19 minimum performance requirements. Where the work flow controller 70 has established multiple
20 branches, each such branch may include minimum performance requirements. The branch-specific
21 performance requirements may be established by the work flow controller 70 based on overall
22 performance requirements for the job ticket 61. A processor 80 that bids on a particular branch
23 may be rejected by the bidding service 90 if the processor 80 fails to meet the minimum
24 performance requirements.

25 If the client 31 does not specify any minimum performance requirements, the bidding
26 service 90 may apply a standard set of criteria (e.g., an industry standard). In addition, the bid
27 must satisfy any requirements for producing output resources. In this way, bids that are made in
28 error, or that would otherwise likely be rejected, can be screened out. For example, a bid for
29 printing inside pages of the brochure may indicate a one year completion date. Such a bid may be
30 rejected, even in the absence of any specified performance requirements from the client 31.

1 In addition to submitting performance requirements, the client 31 may specify an evaluation
2 algorithm for evaluating bids. For example, the client 31 may specify that cost is to be weighted
3 twice as much as any other performance requirement.

4 In the absence of a client-specified evaluation algorithm, the bidding service 90 may apply
5 a standard evaluation algorithm in order to rank bids for each branch in the work flow. The
6 evaluation algorithm may apply weighting criteria, or may apply a default rule. For example, bids
7 may be ranked based on a maximum score, where points are awarded for cost estimates below
8 a maximum and for completion times below a maximum. Once the evaluation algorithm has been
9 applied, the bidding service 90 ranks the bids for each branch. If only one processor 80 survives
10 the process, that processor 80 may be automatically selected and assigned to the branch. If
11 multiple processors 80_i survive, the bidding service 90 may provide a list of such processors 80_i
12 to the work flow controller 70, which will then select the processors 80_i to be assigned to the
13 branches. Alternatively, the list may be provided to the client 31, and the client 31 may select the
14 processor(s) 80_i to complete the tasks defined in the work flow.

15 The work flow controller 70 may associate winning bids with corresponding branches, and
16 may store the bid information with the job ticket 61. The stored bid information may include
17 identification information that allows the authorization server 94 and the authentication server 92
18 to permit access to job ticket branches or to the entire job ticket 61. Because the bid information
19 is stored with the job ticket 61, a processor 80 may access those branches for which the processor
20 80 is authorized access without having to communicate directly with the job ticket service 61. This
21 feature allows the job ticket 60 to be passed from one processor 80 to another processor 80,
22 which improves processing time and efficiency.

23 In an embodiment, the work flow controller 70 accesses control data of the job ticket 61
24 to determine which processor(s) 80_i should be assigned to the specific task identified in the job
25 ticket. The work flow controller 70 may also identify which of the processors 80_i would be able
26 to meet the criteria specified in the control data, and may provide a list of such processors 80_i
27 to the client through the front end service 30. The client 31 may then select a processor(s) 80_i from
28 the list.

29 The job ticket service may, in an alternative embodiment, be implemented in whole or in
30 part by storing program instructions on a computer-readable medium, such as a CD-ROM, for

example. A computer processor may then implement method steps to provide the job ticket service by reading and executing the program instructions.

Figure 5A illustrates an exemplary job ticket 61. The job ticket 61 may include two parts. A first part includes a framework 62 and an optional client extension 64. The framework 62 includes information, files and programming necessary to control tasks defined in the job ticket 61. The client extension 64 may include information related to a specific client (machine) and to a user of the machine. A second part includes a security module 67 that protects the job ticket 61 from unauthorized access.

The framework 62 may include a job identification (ID) 63, a service ID 65, a task section 68, and control data 69. The job ID 63 includes a reference to a specific job, or content 51 that is stored in the job store 50. The job ID 63 also includes a reference to a particular job store 50 that is used to store the content 51. An entity that acquires a reference to the job ticket 61 can use the job ID 63 to access the corresponding content 51. Thus, the network 20 shown in Figure 3 may include multiple job stores 50, and the job ID 63 may be used to correlate the job ticket 61 to a specific job store 50. The service ID 65 identifies a specific job ticket service 60 that stores the job ticket 61. For example, the network 20 may include multiple job ticket services 60 (not shown in Figure 3). The service ID 65 is used to correlate the job ticket 61 to the appropriate job ticket service 60.

The tasks section 68 (Figure 5B) may include branch definitions, and other information needed to control completion of the branches. The tasks section 68 may be structural so that each branch or node in a node tree is represented by one or more branches 66, in the tasks section 68. In this embodiment, each node in the node tree (e.g., the mode tree 10 of Figure 2) can have associated with the node, the description 95, resources 96, lock/unlock flag 97, and security features 99. In this way, the job ticket 61 reflects a hierarchical database structure. The control data 69 includes the specific instructions, parameters, and criteria for completing the task identified by the job ticket 61. The control data 69 may also include specific data required to complete tasks defined in the job ticket 61. The control data 69 may also be associated with each node in a node tree.

Figure 5C illustrates an embodiment of the control data 69. The control data 69 includes a client address, which may be a machine address, such as an Internet protocol (IP) address. An expiration date/time segment may be used to terminate active status of the ticket 61. Once terminated, the ticket may be deleted from the job ticket service, and the corresponding content 51 may be de-referenced. That is, the contents may no longer be referenced by a specific job ticket 61. This feature may help eliminate stale data, and free up resources for other job requests 32 (see Figure 7). Finally, the control data 69 may include specific performance requirements, such as cost an delivery, warranty, required materials, price reductions based on quantity, and other requirements, for example.

The use of job tickets as XML objects allows clients to define databases, and to store data through the job ticket service 60 and the job store 50. The databases may be used to hold contact lists, addresses, and other personal data. The databases may also be used to store any other generic data. The databases could then be used in conjunction with a variety of e-services provided by the processors 80. For example, an e-mail processor 80 that provides e-mail services may be used in conjunction with a personal contact list to send e-mail messages, transfer electronic files, or to establish a chat room. The e-mail processor 80 may access the contact list at predefined intervals to send e-mail messages to a select group of e-mail addressees. Furthermore, because the service center 40 provides a single portal to processors 80 that are coupled to the communications network 35, the client 31 need not have any knowledge of the database structure, or the processing requirements of the processors 80.

In the specific application of the generic XML database to an e-mail service, the client 31 may have established, as a generic database, a list of e-mail contacts. The contacts database may then be stored in the job store 50 as a content file 51. A corresponding job ticket 61 may be stored at the job ticket service 61. The job ticket 61 includes control data needed to send and receive e-mail through the service center 40. Furthermore, the job ticket 61 serves as a pointer to data in the content file 51. In particular, the job ticket 61 may store XML data that is related to other data stored in the content file 51.

Alternatively, the job ticket 61 may store the contacts data. This alternative takes advantage of the fact that the job ticket 61 includes a vocabulary that can be extended to include the contact data, and that the vocabulary can be further extended to include properties for each

1 contact in the contact data. For example, the job ticket 61 may specify that a contact is a business
2 contact or a personal contact. Other properties may also be included, such as whether the contacts
3 in the contact database use mobile phones, land line phones, facsimile machines, and e-mail
4 addresses.

5 The use of the job ticket 61 also allows for parsing, searching and updating the contacts
6 database. For example, the client 31 may desire to search the contacts database for phone
7 numbers for all persons whose first name is Joe. This search functionality is included in the job
8 ticket 61, and allows the job ticket service 60 to provide the client with a list of phone numbers for
9 all entries in the contacts database where the person's first name is Joe. That is, the contacts
10 database includes entries having the property of Joe, and the job ticket service is able to search the
11 contacts database for this property, and to return a list of those entries to the client 31.

12 The properties function of the job ticket 61 also allows the job ticket service 60 to control
13 specific tasks desired by the client 31, or to indicate to the client that a desired task cannot be
14 completed. Staying with the example of the contacts database, the client 31 may desire to send
15 a facsimile transmission to all entries in the contact list that have a specific zip code. The job ticket
16 service 60 can search the contacts database by properties, looking for zip code. The job ticket
17 service 60 can also search the contacts database to determine if any entry does not have a facsimile
18 machine. For those entries that do not have a facsimile machine, the job ticket service 60 can
19 originate a message to send back to the client 31, informing the client 31 that the facsimile
20 transmission was undeliverable. Using this functionality, the client 31 need not know anything about
21 the intended recipients of the facsimile transmission.

22 Returning to the example of an e-mail service, at the client 31, an e-mail application may
23 be launched in order to send an e-mail message, using the Internet, to one or more contacts in the
24 contact database. However, the client 31 need not subscribe to any one Internet service provider.
25 Instead, the service center 40 determines which processor 80 best suits the client's needs for
26 sending the e-mail message. That is, the service center 40 may select a e-mail service provider (a
27 processor 80) to send the e-mail message to a chosen destination address. Furthermore, the
28 service center 40 may determine, based on information maintained in the contact database (i.e., the
29 content 51 in the job store 50); which delivery options are desired by a user at the destination
30 address. For example, the destination address user may desire that all e-mail messages be sent to

09873192 060501
"05090" 267E2B60

1 an e-mail box, or that an alert be provided whenever an e-mail message is sent. These delivery
2 features may be stored in the contact database in the job ticket 61. Alternatively, the delivery
3 features may be stored in a separate database (content file 51) in the job store 50, and the service
4 center may retrieve information from this separate database when determining how to deliver the
5 e-mail message. Specifically, the separate database may include a variety of users, along with the
6 user's Internet address. By comparing the Internet address provided with the out going e-mail to
7 the Internet addresses in the separate database, the service center 40 can determine desired
8 delivery options of the addressee. This process for determining delivery options is transparent to
9 the client 31 that originated the e-mail message. All that the client 31 need know is the contact
10 information (e.g., the Internet address or a contact name).

11 The client 31 may use the job ticket service 60 to specify a number of performance features
12 related to the e-mail service. For example, the client 31 may want the service center 40 to attempt
13 a specified number of delivery attempts, and if delivery does not occur, to send a return message
14 to the client 31 indicating non-delivery of the e-mail message.

15 In another application, the service center 40 may operate as a focal point for a client's
16 messaging systems. The service center may receive and store messages in the job store 50, or
17 within a job ticket 61. Preferably, a separate content file 51 or job ticket 61 is established for each
18 client 31 having an account on the service center 40 so that all of the messages for a single client
19 31 will be stored in the same directory.

20 The network 20 shown in Figure 4 may include the necessary communications interfaces
21 to support voice, e-mail, and facsimile transmissions. When a call arrives at the service center 40,
22 a central processor (not shown in Figure 4) may read an incoming address signal from the
23 originating user, and use the address to store the message in the appropriate job ticket 61.

24 In addition to handling incoming calls and storing the messages, the service center may
25 coordinate a response from the client 31. For example, the client may request a specific response
26 (e.g., and out of office notice) to an e-mail message from a user. The responses may be stored as
27 part of the client's content file 51. The job ticket service 60 may access the content file 51, and
28 download an appropriate response message. The job ticket service 60 may then use a processor
29 80 to deliver the response back to the originating user.

1 The job ticket service 60 may also include software to transform incoming messages (voice,
2 facsimile) into XML files for storage in the content file 51.

3 The job ticket 61 or the content file 51 may contain the data indicating the preferences of
4 the client 31. Thus, for example, when a facsimile message in the TIFF/F format is retrieved by the
5 service center 40, the job ticket service 60 may ascertain from the data in job ticket 61 the
6 preferred option of displaying the facsimile message and would generate the appropriate HTML
7 files.

8 The service center 40 may be connected to a paging system processor 80. Upon the
9 arrival of a new message, in addition to sending an e-mail message to the client's mailbox, the
10 service center 40 may also activate the paging system processor 80 so that a pager (not shown)
11 would be activated. In this manner, the client 31 could receive almost instantaneous notification that
12 a message has arrived.

13 The service center 40 may use the data in the job ticket 61 to provide other information
14 and alerts to the client. For example, the service center 40 may send a message to the client to
15 indicate upcoming appointments, and special dates, such as birthdays. The client 31 may designate
16 in the control data section of the job ticket 61 that birthday greeting be automatically sent to specific
17 individuals whose information is stored in the job ticket 61.

18 The service center 40, when supporting client communications, may operate according an
19 asynchronous manner. In an asynchronous mode of operation, information requested by the client
20 31 may not be available and may have to be generated after the job request 32. The asynchronous
21 mode of operation is preferred since fewer files are generated, thereby reducing the required
22 amount of storage. Because the information requested by a client 31 may not be available, some
23 anchors cannot specify the filename, such as "2.html," but will instead contain a command for the
24 file. For instance, an anchor may be defined as<AHREF=/faxweb/clients/2496801/
25 viewpage.cgi?FAX.sub.--NUM=1&PAGE=1&VIEW.sub.--MODE=FULL> for causing a CGI
26 to run a viewpage program so that page 1 of facsimile message 1 will be displayed in a full size
27 image. The CGI will generate the requested information when the information has not been
28 generated, otherwise the CGI will retrieve the information and relay the information for transmission
29 to the client 31.

30 The service center 40 can reliably receive voice, facsimile, and data messages for a plurality

1 as an archive for a particular client 31 or group of clients 31. As described above, the service
2 center 40 may maintain a list of all messages for a particular client 31, which is displayed to the
3 client 31 when the client 31 accesses a mailbox. The service center 40 may store all messages,
4 whether they are voice, facsimile, or data, for a client 31 in the database indefinitely. The service
5 center 40 may therefore be relied upon by a client 31 to establish the authenticity of a message and
6 the existence or absence of a particular message. Through the service center 40, a client 31 can
7 therefore maintain an accurate record of all received email messages, facsimile messages, and data
8 transfers.

9 In addition to serving as a file depository, the service center 40 may also function as a
10 document management tool. When the service center 40 receives a message, the service center
11 40 updates a database (content file 51) with information on the message. This information includes
12 the type of message, whether it is a facsimile message, voice message, or data message, the time
13 and date at which the message was received, the size of the file, such as in bytes, the telephone
14 number of the caller leaving the message, as well as other information, such as the number of pages
15 of a facsimile message. Because the address (e.g., a telephone number or e-mail address) called
16 is unique for each client 31, the information also includes the intended recipient of the message.

17 As noted above, the job ticket 61, in conjunction with other components of the service
18 center 40, may also be used to create a persistent, generic object-based data structure, such as an
19 XML database. An example of the use of a job ticket 61 for this purpose is illustrated in Figure
20 5D. The job ticket 61 includes a contacts list 84, which may be in the form of an XML database,
21 or some other generic database. The contacts list 84 may include a structure with entries for
22 business 85 and personal 86 use. The business 85 and personal 86 contacts structures may include
23 entries of individuals 87, as shown. Each of the entries 87 may include specific properties, as
24 defined above. In addition, or alternatively, each of the entries 87 may include links to other
25 databases that provide additional information and properties about the individual.

26 While the use of the job ticket 61 as a XML database has generally been described with
27 reference to an e-mail and messaging service, the job ticket 61 is not so limited. Any data that is
28 capable of being stored in a database may be accesses and controlled using the job ticket 61.

29 The features described above, and shown in Figures 5A-5D may be replicated in another
30 embodiment of a job ticket 61 in which all data related to a specific node or branch is located with

that node or branch. Using the example node-tree 10 shown in Figure 2, each node (branch) may include detailed information, and features such as resources, authorized processors 80_i, lock/unlocked flag, bid information, branch description, and other information.

Figure 6 is a diagram of functions of the job ticket service 60. The primary functions of the job ticket service 60 are to store 73 the job tickets 61_i and to provide access 75 to the job tickets 61_i to users such as the client 31 and to the processors 80_i. To accomplish these storage and access functions, the job ticket service 60 may create a job ticket reference 72 and a job resource reference 74. The job ticket service 60 also controls job content access 76, updates 77 the job tickets 61_i as processes are completed and reported by the processors 80_i, completes the job tickets 61_i and reports 78 when all processes are completed for a specific job ticket 61_i, and provides an approval process 79 to allow a client 31 to approve completion of the tasks designated in the job ticket 61.

The job ticket reference 72 includes a specific reference to a corresponding job ticket 61_i. The job ticket reference 72 may be used by the job ticket service 60 to allow one or processors 80_i and clients 31_i to access the job ticket 61. That is, instead of passing the job ticket 61 to a processor 80, the job ticket service 60 passes the job ticket reference 72. With the job ticket reference 72, the processor 80 may access all or a part of a job ticket 61 so that the processor 80 may complete one or more processes. Unlike conventional job ticket services, the job ticket service 60 retains the job ticket in storage 73, and only permits users (clients 31 and processors 80) to access the job ticket 61. This feature allows multiple processors 80 to simultaneously complete processes for the specific job request 32 related to the job ticket 61.

The job ticket service 60 may also create a resources reference 74, and may provide the resources reference 74 to the processors 80 and the clients 31 in a manner similar to that of the job ticket reference 72. As noted above with the description accompanying Figure 2, the resources may include physical devices and materials, and may include digital files. Use of the resources reference 74 may simplify data included in the job ticket 61.

Alternatively, information contained in the resources reference 74 may be included in the job ticket 61, or may be included in other files accessed by the clients 31_i and the processors 80_i.

Figure 7 is a diagram showing operation of selected functions of the job ticket service 60. As shown in Figure 7, the job ticket service 60 includes a job ticket 61_i, which may be a

programming object such as that represented in Figure 2, and described above. The job ticket 61_i is shown supplied to the job ticket service 60 by the client 31_i. The client 31_i may be a networked computer or similar device that is capable of transmitting the digital information representing the job ticket 61_i to the job ticket service 60. To ensure the job ticket 61_i arrives at the job ticket service 60, the job ticket 61_i may contain a reference to the job ticket service 60, such as the service ID 65 illustrated in Figure 5A. The service ID 65 may include a network address of the job ticket service 60. For example, the service ID 65 may include a universal resource locator (URL) if the job ticket service 60 is an Internet web site.

Also shown in Figure 7 are client 31₂ and processors 80₁ - 80_N. The processors 80₁ - 80_N may include networked resources such as networked printers, electronic-commerce entities, such as Internet web sites, and "brick and mortar" entities, such local print shops that are coupled to the job ticket service 60 using the service bus 41.

The client 31 generates a job request 32 (content 51 and job ticket data). Using the front end service 30 (not shown in Figure 7) and the service bus 41, the client 31_i sends the job ticket data to the job ticket service 60 and the content 51 (not shown in Figure 7) to the job store 50. The job ticket service 60 may pass the job ticket data to the work flow controller 70, which will create a job ticket 61. The content 51_i and the job ticket 61_i are related by the job ID 63. The job ID 63 also includes an identification of the job store 50, and a location within the job store 50 in which the content 51_i is stored. In an alternate embodiment, the content 51_i may be stored at the client 31_i, and may then be accessed by other users through the service bus 41 and the front end service 30.

The job ticket 61_i specifies processes that must be completed to finish the job request 32. As noted above, Figure 2 illustrates processes required to print a brochure, including the inside pages and the cover. More than one processor 80_i may be required to complete such a job request, or to complete the job request in the most cost-efficient and/or timely manner. The work flow controller 70 (not shown in Figure 7) can determine which of the processors 80₁-80_N should complete a specific process, and, if necessary, the order in which such processes should be completed. The work flow controller 70 may poll the various processors 80_i to determine which may be used to complete the job request. The work flow controller 70 may then notify selected processors 80_i that a job request has been registered with the job ticket service 60.

1 For each job ticket 61_i received, the job ticket service 60 creates a reference 72_i to the
2 job ticket 61_i. The processor 80₁ may request access to the job ticket 61 in order to complete one
3 or more processes. In response, the job ticket service 60 provides the processor 80₁ with the job
4 ticket reference 72₁. The job ticket reference 72₁ is then used as an index to the job ticket 61₁.
5 The job ticket reference 72₁ may also be provided to other processors, such as the processor 80₂,
6 and to other clients, such as the client 31₂. The processor 80₂ and the client 31₂ may then access
7 the job ticket 61₁ at the same time as the processor 80₁ accesses the job ticket 61₁. This
8 simultaneous access allows different processes to be completed in parallel. In the example
9 illustrated in Figure 2, the processor 80₁ may complete some or all the processes for the inside
10 pages, and the processor 80₂ may complete the processes for the cover.

11 Figure 8 is a block diagram illustrating an example application of the control features of the
12 job ticket service 60. The job ticket 61₁ is referenced to the job content 51₁ by the job ticket ID
13 63, and information related to the job ticket 61₁ and the job content 51₁ is passed over the service
14 bus 41. The processors 80₁ can access the job content 51₁ and the job ticket 61₁ using the service
15 bus 41. In the illustrated example, the job ticket 61₁ refers to a job request 32 to print a brochure
16 using the processes outlined in Figure 2. The processor 80₁ is designated by the work flow
17 processor 70 to produce the inside pages of the brochure and the processor 80₂ is designated to
18 produce the brochure cover. The processor 80₁ passes a job ticket access request to the job
19 ticket service 60. The access request may include security information that allows the processor
20 80₁ to access the job ticket 61₁ and the corresponding content 51₁ or job. In response, the job
21 ticket service 60 provides a job ticket reference 62₁ that is used by the processor 80₁ to access
22 the job ticket 61₁. The processor 80₁ may use information in the job ticket 61₁ to access the
23 content 51₁ stored in the job store 50. Since the processor 80₁ will produce only the inside pages,
24 the processor 80₁ will not need access to all the information contained in the job ticket 61₁.
25 Furthermore, because the job ticket 61₁ remains in the job ticket service 60, other entities, such
26 as the processor 80₂, may continue to access the job ticket.

27 As the processor 80₁ completes various processes, the processor 80₁ may update the
28 content 51₁ and the job ticket 61₁. Thus, the job ticket 61₁ may reflect the latest status of the job
29 request 32. The status reports may indicate when a node in the node tree 10 is completed, when
30 an interim deadline is completed, when another processor may be used to complete a process, and

1 when all processing is complete. The status report may be included in a digital file that is used by
 2 the work flow controller 70, for example. The status report may also be included in a human-
 3 readable format, such as a pop-up window on a computer display screen. The processor 80₁ may
 4 receive the job ticket reference 72₁, and may complete all scheduled processes, returning the job
 5 ticket reference 72₁ to the job ticket service 60. The processor 80₁ may also send a copy of the
 6 job ticket reference 72₁ to the processor 80₂, so that the processor 80₂ may access the job ticket
 7 61₁, and the content 51₁ and produce the brochure cover.

8 Figure 9 is a flowchart illustrating an operation 100 of the job ticket service 60. The
 9 operation 100 is based on completing the inside pages nodes shown in Figure 2. The operation
 10 100 may be at least partly under the control of the work flow controller 70, or some equivalent
 11 device. The operation 100 assumes that a job request 32 (job ticket data and content) have been
 12 passed to the service center 40, and that a job ticket service 60 has been created. The operation
 13 100 begins at start block 101. In review and assign processors block 105, the work flow
 14 controller 70 determines which processors 80_i are able and available to complete the job. The
 15 work flow controller 70, or the optional bidding service 90 may use polling or bidding features to
 16 make the determination. If more than one processor 80_i is available, and can satisfy the
 17 requirements of the job ticket 61, the work flow controller 70 may assign one specific processor
 18 80 to the job. Alternatively, the work flow controller 70 may provide a list of processors 80_i to
 19 the client 31, and allow the client 31 to select one or more processors 80_i.

20 In request job ticket block 110, a processor 80, having been authorized access to a job
 21 ticket 61, sends an access request to the job ticket service 60 using the service bus 41. In block
 22 115, the job ticket service verifies that the processor 80 may access the job ticket 61. Access may
 23 be controlled by a password, an identification, and a public key/private key security system, for
 24 example. In block 115, if the processor 80 is denied access, an error signal may be sent to the
 25 processor and/or the client 31, block 120.

26 In block 115, if access is authorized, the job ticket service 60 provides the processor 80
 27 with a copy of the job ticket reference 72 corresponding to the job ticket 61, block 125. The job
 28 ticket reference 72 allows the processor 80 to access the job ticket at any time. By accessing the
 29 job ticket 61 at any time, the processor 80 is able to view an updated version of the job ticket 61
 30 as changes are made to the job ticket 61 by other entities, including other processors 80.

105090 23E7860
09873192 060501

1 In block 130, the job store 50 provides access to the job content 51 that is referenced by
2 the job ticket 61. Only that part of the content 51 that may be needed by the processor 80 may
3 be supplied by the job store 50. For example, if the processor 80 is only to generate the inside
4 pages of the brochure, the job store 50 may not provide access to the content required to produce
5 the brochure cover. After receiving the job ticket reference 72 and the content 51, the processor
6 80 may perform one or more tasks using input resources to produce an interim or final output
7 resource. With completion of each node in the node tree 10, the processor 80 may provide an
8 input to the job ticket service 60 to allow modification of the job ticket 61, block 135. If the
9 processor 80 completes all required processes, the processor 80 may provide a final status report
10 to the job ticket service 60, block 140, along with any final modifications to the job ticket 61.

11 In block 145, the job ticket service 60 and the work flow controller 70 determine if any
12 additional tasking may be required. If additional tasks are required, the work flow controller 70
13 will ensure the appropriate processors 80 are assigned, and the operation returns to block 110.
14 If no additional processes are required, the operation moves to block 150 and ends.

15 Figure 10 is a flowchart illustrating the routine 105 for developing a work flow and assigning
16 processors to the work flow. The process starts in block 200. In block 205, the service center
17 40 receives a job request 32. The job request 32 may specify performance requirements,
18 resources, and other parameters, and may include content 51, or a link to the content 51. In block
19 210, the work flow controller 70 defines a work flow to accomplish the tasks specified in the job
20 request 32. The work flow may be represented by a node tree, such as the node tree 10 shown
21 in Figure 2.

22 In block 230, the work flow controller 70 generates a job ticket 61 using the information
23 provided by the job request 32, the work flow generated in block 210, and an appropriate job
24 ticket template. The job ticket 61 is then stored in the job ticket service 60. Any content 51 may
25 be stored in the job store 50.

26 The work flow controller 70 or the job ticket service 60 may create a job ticket notice, or
27 other object, and may post the notice, block 250, at the service center 40 so that outside entities
28 (e.g., the processors 80) may acquire sufficient information to bid on completion of the job ticket
29 61, or a branch 66 of the job ticket 61. In an alternative embodiment, the job ticket 61 may be
30 posted at the service center 40. If the job ticket 61 is posted, the job ticket 61 may include

1 mechanism to limit access to the job ticket or to limit access to certain portions of the job ticket 61.
2 For example, the client extension 64 may not be accessible to the processors 80.

3 In block 270, the service center 40 receives bids from specific processors 80 and in block
4 290, the service center 40 evaluates the bids. In block 295, the service center 40 determines if the
5 client 31 submitting the job request 32 intends to select the winning bid(s), or if the service center
6 40 makes the selection. If the client is to make the selections, in block 300, the service center 40
7 provides the bid information to the client 31. Then, in block 305, the service center 40 receives
8 the selections from the client 31. If the service center 40 is to make the selections, in block 310,
9 the service center 40 selects the winning bid(s). In block 315, the service center notifies the
10 winning processors. The service center may also store the bid information with the corresponding
11 job ticket 61. In block 320, the routine 105 ends.

12 Figure 11 is a flowchart illustrating the sub-routine 210 for defining a work flow. The sub-
13 routine 210 starts in block 350. In block 355, the work flow controller 70 determines if the work
14 flow will contain multiple branches. If the work flow will contain multiple branches, the work flow
15 controller 70 defines the branches, block 360. In block 365, the work flow controller 70 selects
16 a branch for which resources and processes are to be defined. In block 370, the work flow
17 controller 70 defines input resources for a first process, or node. In block 375, the work flow
18 controller 70 defines the tasks to be completed for the first process. In block 380, the work flow
19 controller 70 determines the output resources of the first process. In block 385, the work flow
20 controller 70 determines if another process is required for the work flow or branch. In no
21 additional processes are required, the work flow controller 70 determines if another branch is to
22 be defined, block 390. If another branch is to be defined, the work flow controller 70 selects
23 another branch, block 365, and the sub-routine 210 continues. If another branch is not to be
24 defined, the sub-routine 210 ends, block 395. The results of the work flow definition may be
25 incorporated into the job ticket 61 (see Figure 10, block 230).

26 Figure 12 is a flow chart illustrating the sub-routine 250 of posting a job ticket notice or job
27 ticket. The sub-routine 250 starts in block 400. In block 405, the work flow controller 70
28 determines if the work flow associated with the job ticket 61 includes multiple branches. If the
29 work flow does not include multiple branches, the work flow controller posts the job ticket notice
30 listing the single branch, block 410. If the work flow includes multiple branches, the work flow

1 controller 70 posts the job ticket notice with multiple branches, block 420. The sub-routine 250
 2 then ends.

3 Figure 13 is a flow chart illustrating the sub-routine 290 for evaluating bids. The sub-
 4 routine starts in block 440. In block 445, the bidding service 90 selects a first bid for analysis. In
 5 block 450, the bidding service 90 determines if the client 31 has supplied any evaluation criteria or
 6 requirements. If the client has not supplied evaluation requirements, the bidding service 90
 7 compares the selected bid to a set of standard, minimum performance requirements, which may be
 8 industry-standard requirements, block 455. In block 460, the bidding service 90 determines if the
 9 bid meets the minimum performance requirements. If the bid does not meet the minimum
 10 performance requirements, the bid is rejected, block 475. If the bid is rejected, the bidding service
 11 90 determines if additional bids were submitted, block 495. If additional bids were submitted, the
 12 bidding processor 90 returns to block 445 and selects the next bid for evaluation.

13 If in block 450, the client 31 has supplied performance requirements, the bidding service
 14 90 compares the selected bid to the client-supplied performance requirements, block 465. In
 15 block 470, the bidding service 90 determines if the selected bid meets the minimum criteria of the
 16 client-supplied performance requirements. If the minimum criteria are not met, the bidding service
 17 90 rejects the bid, block 475.

18 In blocks 470 and 460, if the minimum criteria are met, the bidding service 90 determines
 19 if the client 31 has supplied an evaluation algorithm. If the client 31 has not supplied an evaluation
 20 algorithm, the bidding service applies a standard evaluation algorithm, which may be an industry-
 21 standard algorithm, block 485. If the client has supplied an evaluation algorithm, the bidding service
 22 90 applies the client-supplied evaluation algorithm, block 490. The bidding service 90 may then
 23 store the results of the algorithm pending evaluation of all bids.

24 In block 495, the bidding service 90 determines if any bids remain to be evaluated. If
 25 additional bids remain, the sub-routine 290 returns to block 445, and the bidding service selects
 26 the next bid for evaluation. In block 495, if no additional bids remain for evaluation, the bidding
 27 service 90 ranks the bids, block 500. The sub-routine 290 then ends, block 505.

28 Figure 14 is a flowchart illustrating the routine 130 for providing access to a job ticket 61.
 29 The routine 130 begins in block 510. In block 515, the job ticket service 60 receives a job ticket
 30 reference 72 from a processor 80, and retrieves the corresponding job ticket 61, block 520.

In block 525, the job ticket service 60 compares the processor identification to processors listed in the job ticket 61 or branches 66 of the job ticket 61. The job ticket service 60 determines if the selected branches 66 are locked, block 530. If the selected branches 66 are not locked, the job ticket service 60 copies the selected branches 66 to the processor 80, block 535. In block 550, the job ticket service 60 then determines if the selected branches 66 require locking. If the selected branches do not require locking, the routine 130 ends, block 560. If the selected branches 66 require locking, the job ticket service 60 locks the selected branches 66, block 555. The routine 130 then ends, block 560.

In block 530, if the selected branches 66 are locked, the job ticket service 60 determines if the processor 80 intends to modify information in the selected branches 66, block 540. If the processor 80 will not modify the selected branches 66, the job ticket service 60 may provide an error message, block 545. If the selected branches 66 will be modified, the job ticket service 60 may unlock the selected branches 66, block 547.

Figure 15 is a flow diagram of a method for allowing access to a job ticket 61. The method may execute as part of the routine 115 shown in Figure 9. The method starts with block 600. In block 605, the authentication server 94 receives authentication information from a processor 80 and retrieves a job ticket 61 corresponding to a job ticket reference 72 possessed by the processor 80. At this stage of the process, the job ticket 61 (excluding the public key signature field 67) contains two information fields, the framework 62 and the client extension 64. The framework 62 contains information such as the service ID, client IP address, expiration date and time, and processor authorization, as previously described. The client extension 64 contains information such as credit card number and zip code, also previously described. The information in the job ticket 61 (excluding the public key signature field 67) is then, for example, optionally hashed using, for example, MD5 protocol, and encrypted with a public key encryption system, block 610, generating a hash number, block 615. Other hashing or encryption techniques may also be used. The hash number is representative of the specific information contained in the job ticket 61. The hash number generated in block 615 is then encrypted using a standard public key encryption system, block 620. Encrypting the hash number with a private key prevents any user without knowledge of the public key from modifying the job information. In block 625, the job ticket 61 and the encrypted hash number are concatenated to generate the completed job ticket 61. Hence, the completed job ticket

1 devices on which a finite state machine capable of implementing the flowcharts shown in Figures
2 9-16 can be used as the service center 40, or its sub-components.

3 The terms and descriptions used herein are set forth by way of illustration only and are not
4 meant as limitations. Those skilled in the art will recognize that many variations are possible within
5 the spirit and scope of the invention as defined in the following claims, and their equivalents, in
6 which all terms are to be understood in their broadest possible sense unless otherwise indicated.

09873193 060501
"05090" 26TE/860